arXiv:1302.5445v1 [cs.DS] 21 Feb 2013

Thrifty Algorithms for Multistage Robust Optimization

Anupam Gupta* Viswanath Nagarajan† Vijay V. Vazirani‡

February 25, 2013

Abstract

We consider a class of multi-stage robust covering problems, where additional information is revealed about the problem instance in each stage, but the cost of taking actions increases. The dilemma for the decision-maker is whether to wait for additional information and risk the inflation, or to take early actions to hedge against rising costs. We study the "k-robust" uncertainty model: in each stage $i = 0, 1, \ldots, T$, the algorithm is shown some subset of size k_i that completely contains the eventual demands to be covered; here $k_1 > k_2 > \cdots > k_T$ which ensures increasing information over time. The goal is to minimize the cost incurred in the *worst-case* possible sequence of revelations.

For the *multistage k-robust set cover* problem, we give an $O(\log m + \log n)$-approximation algorithm, nearly matching the $\Omega\left(\log n + \frac{\log m}{\log\log m}\right)$ hardness of approximation [4] even for $T = 2$ stages. Moreover, our algorithm has a useful "thrifty" property: it takes actions on just two stages. We show similar thrifty algorithms for multi-stage k-robust *Steiner tree*, *Steiner forest*, and *minimum-cut*. For these problems our approximation guarantees are $O(\min\{T, \log n, \log \lambda_{\max}\})$, where $\lambda_{\max}$ is the maximum inflation over all the stages. We conjecture that these problems also admit $O(1)$-approximate thrifty algorithms.

1 Introduction

This paper considers approximation algorithms for a set of multi-stage decision problems. Here, additional information is revealed about the problem instance in each stage, but the cost of taking actions increases. The decision-making algorithm has to decide whether to wait for additional information and risk the rising costs, or to take actions early to hedge against inflation. We consider the model of robust optimization, where we are told what the set of possible information revelations are, and want to minimize the cost incurred in the *worst-case* possible sequence of revelations.

For instance, consider the following multi-stage set cover problem: initially we are given a set system $(U, \mathcal{F})$. Our eventual goal is to cover some subset $A \subseteq U$ of this universe, but we don't know this "scenario" A up-front. All we know is that A can be any subset of U of size at most k. Moreover we know that on each day i, we will be shown some set A_i of size k_i, such that A_i contains the scenario A—these numbers k_i decrease over time, so that we have more information as time progresses, until $\cap_{i=0}^{T} A_i = A$. We can pick sets from $\mathcal{F}$ toward covering A whenever we want, but the costs of sets increase over time (in a specified fashion). Eventually, the sets we pick must cover the final subset A. We want to minimize the worst-case cost

$$\max_{\sigma=\langle A_1, A_2, \ldots, A_T\rangle : |A_t| = k_t \ \forall t} \text{total cost of algorithm on sequence } \sigma \tag{1.1}$$

*Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 15213, USA. Supported in part by NSF award CCF-1016799 and an Alfred P. Sloan Fellowship. Email: anupamg@cs.cmu.edu

†IBM T.J. Watson Research Center, Yorktown Heights, NY 10598, USA. Email: viswanath@us.ibm.com

‡College of Computing, Georgia Institute of Technology, Atlanta, GA 30332-0280. Supported by NSF Grants CCF-0728640 and CCF-0914732, ONR Grant N000140910755, and a Google Research Grant. Email: vazirani@cc.gatech.edu

This is a basic model for multistage robust optimization and requires minimal specification of the uncertainty sets (it only needs the cardinality bounds k_is).

Robust versions of Steiner tree/forest, minimum cut, and other covering problems are similarly defined. This tension between waiting for information vs. the temptation to buy early and beat rising costs arises even in 2-stage decision problems—here we have T stages of decision-making, making this more acute.

A comment on the kind of partial information we are modeling: in our setting we are given progressively more information about events that *will not* happen, and are implicitly encouraged (by the rising prices) to plan prudently for the (up to k) events that will indeed happen. For example, consider a farmer who has a collection of n possible bad events ("high seed prices in the growing season", $\{$"no rains by month i"$\}_{i=1}^{5}$, etc.), and who is trying to guard against up to k of these bad events happening at the end of the planning horizon. Think of k capturing how risk-averse he is; the higher the k, the more events he wants to cover. He can take actions to guard against these bad events (store seed for planting, install irrigation systems, take out insurance, etc.). In this case, it is natural that the information he gets is about the bad events that do not happen.

This should be contrasted with online algorithms, where we are only given events that do happen—namely, demands that need to be immediately and irrevocably covered. This difference means that we cannot directly use the ideas from online competitive analysis, and consequently our techniques are fairly different.[1] A second difference from online competitive analysis is, of course, in the objective function: we guarantee that the cost incurred on the worst-case sequence of revelations is approximately minimized, as opposed to being competitive to the best series of actions for every set of revelations—indeed, the rising prices make it impossible to obtain a guarantee of the latter form in our settings.

Our Results. In this paper, we give the first approximation algorithms for standard covering problems (set cover, Steiner tree and forest, and min-cut) in the model of multi-stage robust optimization with recourse. A feature of our algorithms that make them particularly desirable is that they are "thrifty": *they actually take actions in just two stages*, regardless of the number of stages T. Hence, even if T is polynomially large, our algorithms remain efficient and simple (note that the optimal decision tree has potentially exponential size even for constant T). For example, the set cover algorithm covers some set of "dangerous" elements right off the bat (on day 0), then it waits until a critical day t^* when it covers all the elements that can concievably still like in the final set A. We show that this set-cover algorithm is an $O(\log m + \log n)$-approximation, which almost matches the hardness result of $\Omega(\log n + \frac{\log m}{\log \log m})$ [4] for $T = 2$.

We also give thrifty algorithms for three other covering problems: Steiner tree, Steiner forest, Min-cut—again, these algorithms are easy to describe and to implement, and have the same structure:

> We find a solution in which decisions need to be made *only at two points in time*: we cover a set of dangerous elements in stage 0 (before any additional information is received), and then we cover all surviving elements at stage t^*, where $t^* = \text{argmin}_t\, \lambda_t k_t$.

For these problems, the approximation guarantee we can currently prove is no longer a constant, but depends on the number of stages: specifically, the dependence is $O(\min\{T, \log n, \log \lambda_{\max}\})$, where $\lambda_{\max}$ is the maximum inflation factor. While we conjecture this can be improved to a constant, we would like to emphasize that even for T being a constant more than two, previous results and

[1] It would be interesting to consider a model where a combination of positive and negative information is given, i.e., a mix of robust and online algorithms. We leave such extensions as future work.

techniques do not imply the existence of a constant-factor approximation algorithm, let alone the existence of a thrifty algorithm.

The definition of "dangerous" in the above algorithm is, of course, problem dependent: e.g., for set cover these are elements which cost more than Opt/k_{t^*} to cover. In general, this defintion is such that bounding the cost of the elements we cover on day t^* is immediate. And what about the cost we incur on day 0? This forms the technical heart of the proofs, which proceeds by a careful backwards induction over the stages, bounding the cost incurred in covering the dangerous elements that are still uncovered by Opt after j stages. These proofs exploit some net-type properties of the respective covering problems, and extend the results in Gupta et al. [6]. While our algorithms appear similar to those in [6], the proofs require new technical ideas such as the use of non-uniform thresholds in defining "nets" and proving properties about them.

The fact that these multistage problems have near-optimal strategies with this simple structure is quite surprising. One can show that the optimal solution may require decision-making at all stages (we show an example for set cover in Section B.1). It would be interesting to understand this phenomenon further. For problems other than set cover (i.e., those with a performance guarantee depending on T), can we improve the guarantees further, and/or show a tradeoff between the approximation guarantee and the number of stages we act in? These remain interesting directions for future research.

We also observe in Section B.2 that thrifty algorithms perform poorly for multistage robust set-cover even on slight generalizations of the above "k-robust uncertainty sets". In this setting it turns out that any reasonable near-optimal solution must act on all stages. This suggests that the k-robust uncertainty sets studied in this paper are crucial to obtaining good thrifty algorithms.

Related Work. Demand-robust optimization has long been studied in the operations research literature, see eg. the survey article by Bertsimas et al. [2] and references therein. The multistage robust model was studied in Ben-Tal et al. [1]. Most of these works involve only continuous decision variables. On the other hand, the problems considered in this paper involve making discrete decisions.

Approximation algorithms for robust optimization are of more recent vintage: all these algorithms are for two-stage optimization with discrete decision variables. Dhamdhere et al. [3] studied two-stage versions when the scenarios were *explicitly* listed, and gave constant-factor approximations for Steiner tree and facility location, and logarithmic approximations to mincut/multicut problems. Golovin et al. [5] gave $O(1)$-approximations to robust mincut and shortest-paths. Feige et al. [4] considered *implicitly* specified scenarios and introduced the k-robust uncertainty model ("scenarios are all subsets of size k"); they gave an $O(\log m \log n)$-approximation algorithm for 2-stage k-robust set cover using an LP-based approach. Khandekar et al. [8] gave $O(1)$-approximations for 2-stage k-robust Steiner tree, Steiner forest on trees and facility location, using a combinatorial algorithm. Gupta et al. [6] gave a general framework for two-stage k-robust problems, and used it to get better results for set cover, Steiner tree and forest, mincut and multicut. We build substantially on the ideas from [6].

Approximation algorithms for multistage stochastic optimization have been given in [9, 7]; in the stochastic world, we are given a probability distribution over sequences, and consider the average cost instead of the worst-case cost in (1.1). However these algorithms currently only work for a constant number of stages, mainly due to the explosion in the number of potential scenarios. The current paper raises the possibility that for "simple" probability distributions, the techniques developed here may extend to stochastic optimization.

Notation. We use $[T]$ to denote $\{0, \cdots, T\}$, and $\binom{X}{k}$ to denote the collection of all k-subsets of the set X.

2 Multistage Robust Set Cover

In this section, we give an algorithm for multistage robust set cover with approximation ratio $O(\log m + \log n)$; this approximation matches the previous best approximation guarantee for two-stage robust set cover [6]. Moreover, our algorithm has the advantage of picking sets only in two stages. (In Section B.1, we show that an optimal strategy might need to pick sets in all stages.)

The multistage robust set cover problem is specified by a set-system $(U, \mathcal{F})$ with $|U| = n$, set costs $c : \mathcal{F} \to \mathbb{R}_+$, a time horizon T, integer values $n = k_0 \geq k_1 \geq k_2 \geq \cdots \geq k_T$, and inflation parameters $1 = \lambda_0 \leq \lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_T$. Define $A_0 = U$, and $k_0 = |U|$. A *scenario-sequence* $\mathbb{A} = (A_0, A_1, A_2, \ldots, A_T)$ is a sequence of $T+1$ 'scenarios' such that $|A_i| = k_i$ for each $i \in [T]$. Here A_i is the information revealed to the algorithm on day i. The elements in $\cap_{i \leq j} A_i$ are referred to as being *active* on day j.

- On day 0, all elements are deemed active and any set $S \in \mathcal{F}$ may be chosen at the cost $c(S)$.
- On each day $j \geq 1$, the set A_j with k_j elements is revealed to the algorithm, and the active elements are $\cap_{i \leq j} A_i$. The algorithm can now pick any sets, where the cost of picking set $S \in \mathcal{F}$ is $\lambda_j \cdot c(S)$.

Feasibility requires that all the sets picked over all days $j \in [T]$ cover $\cap_{i \leq T} A_i$, the elements that are still active at the end. The goal is to minimize the worst-case cost incurred by the algorithm, the worst-case taken over all possible scenario sequences. Let Opt be this worst-case cost for the best possible algorithm; we will formalize this soon. The main theorem of this section is the following:

Theorem 2.1 *There is an $O(\log m + \log n)$-approximation algorithm for the T-stage k-robust set cover problem.*

The algorithm is easy to state: For any element $e \in U$, let $\mathsf{MinSet}(e)$ denote the minimum cost set in $\mathcal{F}$ that contains e. Define $\tau := \beta \cdot \max_{j \in [T]} \frac{\mathsf{Opt}}{\lambda_j k_j}$ where $\beta := 36 \ln m$ is some parameter. Let $j^* = \operatorname{argmin}_{j \in [T]}(\lambda_j k_j)$. Define the "net" $N := \{e \in U \mid c(\mathsf{MinSet}(e)) \geq \tau\}$. Our algorithm's strategy is the following:

> On day zero, choose sets $\phi_0 :=$ Greedy-Set-Cover(N).
>
> On day j^*, for any yet-uncovered elements e in A_{j^*},
> pick a min-cost set in $\mathcal{F}$ covering e.
>
> On all other days, do nothing.

It is clear that this is a feasible strategy; indeed, all elements that are still active on day j^* are covered on that day. (In fact, it would have sufficed to just cover all the elements in $\cap_{i \leq j^*} A_i$.) Note that this strategy pays nothing on days other than 0 and j^*; we now bound the cost incurred on these two days.

Claim 2.2 *For any scenario-sequence $\mathbb{A}$, the cost on day j^* is at most $\beta \cdot \mathsf{Opt}$.*

Proof: The sets chosen in day j^* on sequence $\mathbb{A}$ are $\{\mathsf{MinSet}(e) \mid e \in A_{j^*} \setminus N\}$, which costs us

$$\lambda_{j^*} \sum_{e \in A_{j^*} \setminus N} c(\mathsf{MinSet}(e)) \leq \lambda_{j^*} |A_{j^*}| \cdot \tau = \lambda_{j^*} k_{j^*} \tau = \beta \cdot \mathsf{Opt}.$$

The first inequality is by the choice of N, the last equality is by τ's definition. ∎

Lemma 2.3 *The cost of covering the net N on day zero is at most $O(\log n) \cdot \mathsf{Opt}$.*

The proof of Lemma 2.3 will occupy the rest of this section; before we do that, note that Claim 2.2 and Lemma 2.3 complete the proof for Theorem 2.1. Note that while the definition of the set N requires us to know Opt, we can just run over polynomially many guesses for Opt and choose the one that minimizes the cost for day zero plus $\tau \cdot k_{j^*}\lambda_{j^*}$ (see [6] for a rigorous argument).

The proof will show that the fractional cost of covering the elements in the net N is at most Opt, and then invoke the integrality gap for the set covering LP. For the fractional cost, the proof is via a careful backwards induction on the number of stages, showing that if we mimic the optimal strategy for the first $j-1$ steps, then the fractional cost of covering the remaining active net elements at stage j is related to a portion of the optimal value as well. This is easy to prove for the stage T, and the claim for stage 0 exactly bounds the cost of fractionally covering the net. To write down the precise induction, we next give some notation and formally define what a strategy is (which will be used in the subsequent sections for the other problems as well), and then proceed with the proof.

Formalizing What a Strategy Means For any collection $\mathcal{G} \subseteq \mathcal{F}$ of sets, let $\mathsf{Cov}(\mathcal{G}) \subseteq U$ denote the elements covered by the sets in $\mathcal{G}$, and let $c(\mathcal{G})$ denote the sum of costs of sets in $\mathcal{G}$. At any day i, the *state of the system* is given by the subsequence $(A_0, A_1, \ldots, A_i)$ seen thus far. Given any scenario sequence $\mathbb{A}$ and $i \in [T]$, we define $\mathbb{A}_i = (A_0, A_1, \ldots, A_i)$ to be the partial scenario sequence for days 0 through i.

A solution is a *strategy* Φ, given by a sequence of maps $(\phi_0, \phi_1, \ldots, \phi_T)$, where each one of these maps ϕ_i maps the state $\mathbb{A}_i$ on day i to a collection of sets that are picked on that day. For any scenario-sequence $\mathbb{A} = (A_1, A_2, \ldots, A_T)$, the strategy Φ does the following:

- On day 0, when all elements in U are active, the sets in ϕ_0 are chosen, and $\mathcal{G}_1 \leftarrow \phi_0$.
- At the beginning of day $i \in \{1, \cdots, T\}$, sets in $\mathcal{G}_i$ have already been chosen; moreover, the elements in $\cap_{j \le i} A_j$ are the active ones. Now, sets in $\phi_i(\mathbb{A}_i)$ are chosen, and hence we set $\mathcal{G}_{i+1} \leftarrow \mathcal{G}_i \cup \phi_i(\mathbb{A}_i)$.

The solution $\Phi = (\phi_i)_i$ is *feasible* if for every scenario-sequence $\mathbb{A} = (A_1, A_2, \ldots, A_T)$, the collection $\mathcal{G}_{T+1}$ of sets chosen at the end of day T covers $\cap_{i \le T} A_i$, i.e. $\mathsf{Cov}(\mathcal{G}_{T+1}) \supseteq \cap_{i \le T} A_i$. The cost of this strategy Φ on a fixed sequence $\mathbb{A}$ is the total effective cost of sets picked:

$$C(\Phi \mid \mathbb{A}) = c(\phi_0) + \textstyle\sum_{i=1}^{T} \lambda_i \cdot c\left(\phi_i(\mathbb{A}_i)\right).$$

The objective in the robust multistage problem is to minimize $\mathsf{RobCov}(\Phi)$, the effective cost under the *worst case* scenario-sequence, namely:

$$\mathsf{RobCov}(\Phi) := \max_{\mathbb{A}} C(\Phi \mid \mathbb{A})$$

The goal is to find a strategy with least cost; for the rest of the section, fix $\Phi^* = \{\phi_i^*\}$ to be such a strategy, and let $\mathsf{Opt} = \mathsf{RobCov}(\Phi^*)$ denote the optimal objective value.

Completing Proof of Lemma 2.3. First, we assume that the inflation factors satisfy $\lambda_{j+1} \ge 12 \cdot \lambda_j$ for all $j \ge 0$. If the instance does not have this property, we can achieve this by merging consecutive days having comparable inflations, and lose a factor of 12 in the approximation ratio. The choice of constant 12 comes from a lemma from [6].

Lemma 2.4 ([6]) *Consider any instance of set cover; let $B \in \mathbb{R}_+$ and $k \in \mathbb{Z}_+$ be values such that*

- *the set of minimum cost covering any element costs $\ge 36 \ln m \cdot \frac{B}{k}$, and*
- *the minimum cost of fractionally covering any k-subset of elements $\le B$.*

Then the minimum cost of fractionally covering **all** *elements is at most $r \cdot B$, for a value $r \le 12$.*

For a partial scenario sequence $\mathbb{A}_i$ on days upto i, we use $\phi_i^*(\mathbb{A}_i)$ to denote the sets chosen *on day* i by the optimal strategy, and $\phi_{\leq i}^*(\mathbb{A}_i)$ to denote the sets chosen on days $\{0, 1, \ldots, i\}$, again by the optimal strategy.

Definition 2.5 *For any $j \in [T]$ and $\mathbb{A}_j = (A_1, \ldots, A_j)$, define*

$$V_j(\mathbb{A}_j) := \max_{\substack{(A_{j+1}, \cdots, A_T) \\ |A_t| = k_t \ \forall t}} \sum_{i=j}^{T} r^{i-j} \cdot c\left(\phi_i^*(\mathbb{A}_i)\right).$$

That is, $V_j(\mathbb{A}_j)$ is the worst-case cost incurred by Φ^* on days $\{j, \ldots, T\}$ conditioned on $\mathbb{A}_j$, under modified inflation factors r^{i-j} for each day $i \in \{j, \ldots, T\}$. We use this definition with r being the constant from Lemma 2.4. Recall that we assumed that $\lambda_i \geq r^i$.

Fact 2.6 *The function $V_0(\cdot)$ takes the empty sequence as its argument, and returns $V_0 = \max_{\mathbb{A}} \sum_{i=0}^{T} r^i \cdot c(\phi_i^*(\mathbb{A}_i)) \leq \max_{\mathbb{A}} \sum_{i=0}^{T} \lambda_i \cdot c(\phi_i^*(\mathbb{A}_i)) = \mathsf{Opt}$,*

For any subset $U' \subseteq U$ and any collection of sets $\mathcal{G} \subseteq \mathcal{F}$, define $\mathsf{LP}(U' \mid \mathcal{G})$ as the minimum cost of *fractionally covering* U', given all the sets in $\mathcal{G}$ at zero cost. Given any sequence $\mathbb{A}$, it will also be useful to define $\widehat{A}_j = \cap_{i \leq j} A_j$ as the active elements on day j. Our main technical lemma is the following:

Lemma 2.7 *For any $j \in [T]$ and partial scenario sequence $\mathbb{A}_j$, we have:*

$$\mathsf{LP}\left(N \cap \widehat{A}_j \mid \phi_{\leq j-1}^*\left(\mathbb{A}_{j-1}\right)\right) \leq V_j(\mathbb{A}_j).$$

In other words, the fractional cost of covering $N \cap \widehat{A}_i$ (the "net" still active in stage j) given sets $\phi_{\leq j-1}^*\left(\mathbb{A}_{j-1}\right)$ for free is at most $V_j(\mathbb{A}_j)$.

Before we prove this, note that for $j = 0$, the lemma implies that $LP(N) \leq V_0 \leq \mathsf{Opt}$. Since the integrality gap for the set cover LP is at most H_n (as witnessed by the greedy algorithm), this implies that the cost on day 0 is at most $O(\log n)\mathsf{Opt}$, which proves Lemma 2.3.

Proof: We induct on $j \in \{0, \cdots, T\}$ with $j = T$ as base case. In this case, we have a complete scenario-sequence $\mathbb{A}_T = \mathbb{A}$, and the feasibility of the optimal strategy implies that $\phi_{\leq T}^*(\mathbb{A}_T)$ completely covers $\widehat{A}_T$. So,

$$\mathsf{LP}\left(\widehat{A}_T \mid \phi_{\leq T-1}^*(\mathbb{A}_{T-1})\right) \leq c\left(\phi_T^*(\mathbb{A}_T)\right) = V_T(\mathbb{A}_T).$$

For the induction step, suppose now $j < T$, and assume the lemma for $j+1$. Here's the roadmap for the proof: we want to bound the fractional cost to cover elements in $N \cap \widehat{A}_j \setminus \mathsf{Cov}(\phi_{\leq j-1}^*(\mathbb{A}_{j-1}))$ since the sets $\phi_{\leq j-1}^*(\mathbb{A}_{j-1})$ are free. Some of these elements are covered by $\phi_j^*(\mathbb{A}_j)$, and we want to calculate the cost of the others—for these we'll use the inductive hypothesis. So given the scenarios $A_1, \ldots, A_j$ until day j, define

$$W_j(\mathbb{A}_j) := \max_{|B_{j+1}| = k_{j+1}} V_{j+1}(A_1, \ldots, A_j, B_{j+1}) \implies V_j(\mathbb{A}_j) = c(\phi_j^*(\mathbb{A}_j)) + r \cdot W_j(\mathbb{A}_j). \tag{2.2}$$

Let us now prove two simple subclaims.

Claim 2.8 $W_j(\mathbb{A}_j) \leq \mathsf{Opt}/\lambda_{j+1}$.

Proof: Suppose that $W_j(\mathbb{A}_j)$ is defined by the sequence $(A_{j+1}, \ldots, A_T)$; i.e. $W_j(\mathbb{A}_j) = \sum_{i=j+1}^{T} r^{i-j-1} \cdot c\left(\phi_i^*(\mathbb{A}_i)\right)$. Then, considering the scenario-sequence $\mathbb{A} = (A_1, \ldots, A_j, A_{j+1}, \ldots, A_T)$, we have:

$$\mathsf{Opt} \geq \sum_{i=0}^{T} \lambda_i \cdot c\left(\phi_i^*(\mathbb{A}_i)\right) \geq \sum_{i=j+1}^{T} \lambda_i \cdot c\left(\phi_i^*(\mathbb{A}_i)\right) \geq \sum_{i=j+1}^{T} \lambda_{j+1}\, r^{i-j-1} \cdot c\left(\phi_i^*(\mathbb{A}_i)\right) = \lambda_{j+1} \cdot W_j(\mathbb{A}_j).$$

The third inequality uses the assumption that $\lambda_{\ell+1} \geq r \cdot \lambda_\ell$ for all days ℓ. ◀

Claim 2.9 *For any A_{j+1} with $|A_{j+1}| = k_{j+1}$, we have*

$$\mathsf{LP}\left(N \cap \widehat{A}_{j+1} \mid \phi^*_{\leq j}(\mathbb{A}_j)\right) \leq W_j(\mathbb{A}_j).$$

Proof: By the induction hypothesis for $j+1$, and $V_{j+1}(\mathbb{A}_{j+1}) \leq W_j(\mathbb{A}_j)$. ◀

Now we are ready to apply Lemma 2.4 to complete the proof of the inductive step.

Claim 2.10 *Consider the set-system $\mathcal{G}$ with elements $N' := N \bigcap \left(\widehat{A}_j \setminus \mathsf{Cov}(\phi^*_{\leq j}(\mathbb{A}_j))\right)$ and the sets $\mathcal{F} \setminus \phi^*_{\leq j}(\mathbb{A}_j)$. The fractional cost of covering N' is at most $r \cdot W_j(\mathbb{A}_j)$.*

Proof: In order to use Lemma 2.4 on this set system, let us verify the two conditions:

1. Since $N' \subseteq N$, the cost of the cheapest set covering any $e \in N'$ is at least $\tau \geq \beta \cdot \frac{\mathsf{Opt}}{\lambda_{j+1}\, k_{j+1}} \geq \beta \cdot \frac{W_j(\mathbb{A}_j)}{k_{j+1}}$ using the definition of the threshold τ, and Claim 2.8; recall $\beta = 36 \ln m$.
2. For every $X \subseteq N'$ with $|X| \leq k_{j+1}$, the minimum cost to fractionally cover X in $\mathcal{G}$ is at most $W_j(\mathbb{A}_j)$. To see this, augment X arbitrarily to form A_{j+1} of size k_{j+1}; now Claim 2.9 applied to A_{j+1} implies that the fractional covering cost for $N \cap \widehat{A}_{j+1} = N \cap \left(A_{j+1} \cap \widehat{A}_j\right)$ in $\mathcal{G}$ is at most $W_j(\mathbb{A}_j)$; since $X \subseteq N' \subseteq N \cap \widehat{A}_j$ and $X \subseteq A_{j+1}$ the covering cost for X in $\mathcal{G}$ is also at most $W_j(\mathbb{A}_j)$.

We now apply Lemma 2.4 on set-system $\mathcal{G}$ with parameters $B := W_j(\mathbb{A}_j)$ and $k = k_{j+1}$ to infer that the minimum cost to fractionally cover N' using sets from $\mathcal{F} \setminus \phi^*_{\leq j}(\mathbb{A}_j)$ is at most $r \cdot W_j(\mathbb{A}_j)$. ◀

To fractionally cover $N \cap \widehat{A}_j$, we can use the fractional solution promised by Claim 2.10, and integrally add the sets $\phi_j^*(\mathbb{A}_j)$. This implies that

$$\mathsf{LP}\left(N \cap \widehat{A}_j \mid \phi^*_{\leq j-1}(\mathbb{A}_{j-1})\right) \leq c(\phi_j^*(\mathbb{A}_j)) + r \cdot W_j(\mathbb{A}_j) = V_j(\mathbb{A}_j),$$

where the last equality follows from (2.2). This completes the induction and proves Lemma 2.7. ■

3 Multistage Robust Minimum Cut

We now turn to the multistage robust min-cut problem, and show:

Theorem 3.1 *There is an $O\left(\min\{T, \log n, \log \lambda_{max}\}\right)$-approximation algorithm for T-stage k-robust minimum cut.*

In this section we prove an $O(T)$-approximation ratio where T is the number of stages; in Appendix A we show that simple scaling arguments can be used to ensure T is at most $\min\{\log n, \log \lambda_{max}\}$, yielding Theorem 3.1. Unlike set cover, the guarantee here depends on the number of stages. Here is the high-level reason for this additional loss: in each stage of an optimal strategy for set cover, any element was either completely covered or left completely uncovered—there was no partial coverage. However in min-cut, the optimal strategy could keep whittling away at the cut for a node in each stage. The main idea to deal with this is to use a stage-dependent definition of "net" in the inductive proof (see Lemma 3.4 for more detail), which in turn results in an $O(T)$ loss.

The input consists of an undirected graph $G = (U, E)$ with edge-costs $c : E \to \mathbb{R}_+$ and root ρ. For any subset $U' \subseteq U$ and subgraph H of G, we denote by $\mathsf{MinCut}_H(U')$ the minimum cost of a cut separating U' from ρ in H. If no graph is specified then it is relative to the original graph G. Recall that a scenario sequence $\mathbb{A} = (A_0, A_1, \ldots, A_T)$ where each $A_i \subseteq U$ and $|A_i| = k_i$, and we denote the partial scenario sequence $(A_0, A_1, \ldots, A_j)$ by $\mathbb{A}_j$.

We will use notation developed in Section 2. Let the optimal strategy be $\Phi^* = \{\phi_j^*\}_{j=0}^T$, where now $\phi_j^*(\mathbb{A}_j)$ maps to a set of edges in G to be cut in stage j. The feasibility constraint is that $\phi_{\le T}^*(\mathbb{A}_T)$ separates the vertices in $\cap_{i \le T} A_i$ from the root ρ. Let the cost of the optimal solution be $\mathsf{Opt} = \mathsf{RobCov}(\Phi^*)$.

Again, the algorithm depends on showing a near-optimal two-stage strategy: define $\tau := \beta \cdot \max_{j \in [T]} \frac{\mathsf{Opt}}{\lambda_j k_j}$, where $\beta = 50$. Let $j^* = \operatorname{argmin}_{j \in [T]}(\lambda_j k_j)$. Let the "net" $N := \{v \in U \mid \mathsf{MinCut}(v) > 2T \cdot \tau\}$. The algorithm is:

On day 0, delete $\phi_0 := \mathsf{MinCut}(N)$ to separate the "net" N from ρ.

On day j^*, for each vertex u in $A_{j^*} \setminus N$, delete a minimum u-ρ cut in G.

On all other days, do nothing.

Again, it is clear that this strategy is feasible: all vertices in $\cap_{i \le T} A_i$ are either separated from the root on day 0, or on day j^*. Moreover, the effective cost of the cut on day j^* is at most $\lambda_{j^*} \cdot 2T\tau \cdot |A_{j^*}| = 2\beta T\, \mathsf{Opt} = O(T) \cdot \mathsf{Opt}$. Hence it suffices to show the following:

Lemma 3.2 *The min-cut separating N from the root ρ costs at most $O(T) \cdot \mathsf{Opt}$.*

Again, the proof is via a careful induction on the stages. Loosely speaking, our induction is based on the following: amongst scenario sequences $\mathbb{A}$ containing any fixed "net" vertex $v \in N$ (i.e. $v \in \cap_{i \le T} A_i$) the optimal strategy must reduce the min-cut of v (in an average sense) by a factor $1/T$ in some stage.

The proof of Lemma 3.2 again depends on a structural lemma proved in [6]:

Lemma 3.3 ([6]) *Consider any instance of minimum cut in an undirected graph with root ρ and terminals X; let $B \in \mathbb{R}_+$ and $k \in \mathbb{Z}_+$ be values such that*

- *the minimum cost cut separating ρ and x costs $\ge 10 \cdot \frac{B}{k}$, for every $x \in X$.*
- *the minimum cost cut separating ρ and L is $\le B$, for every $L \in \binom{X}{k}$.*

Then the minimum cost cut separating ρ and **all** *terminals X is at most $r \cdot B$, for a value $r \le 10$.*

In this section, we assume $\lambda_{j+1} \ge 10 \cdot \lambda_j$ for all $j \in [T]$. Recall the quantity $V_j(\mathbb{A}_j)$ from Definition 2.5:

$$V_j(\mathbb{A}_j) := \max_{\substack{(A_{j+1}, \cdots, A_T) \\ |A_t| = k_t \ \forall t}} \sum_{i=j}^{T} r^{i-j} \cdot c\left(\phi_i^*(\mathbb{A}_i)\right)$$

where $r := 10$ from Lemma 3.3. Since $\lambda_i \geq r^i$, it follows that $V_0 \leq \mathsf{Opt}$. The next lemma is now the backwards induction proof that relates the cost of cutting subsets of the net N to the V_js. This finally bounds the cost of separating the entire net N from ρ in terms of $V_0 \leq \mathsf{Opt}$. Given any $\mathbb{A}_j$, recall that $\widehat{A}_j = \cap_{i \leq j} A_i$.

Lemma 3.4 *For any $j \in [T]$ and partial scenario sequence $\mathbb{A}_j$,*

- *if $H := G \setminus \phi^*_{\leq j-1}(\mathbb{A}_{j-1})$ (the residual graph in OPT's run at the beginning of stage j), and*
- *$N_j := \{v \in \widehat{A}_j \mid \mathsf{MinCut}_H(v) > (2T - j) \cdot \tau\}$ (the "net" elements)*

then $\mathsf{MinCut}_H(N_j) \leq 5T \cdot V_j(\mathbb{A}_j)$.

Before we prove the lemma, note that when we set $j = 0$ the lemma claims that in G, the min-cut separating $N_0 = N$ from ρ costs at most $5T \cdot V_0 \leq O(T)\,\mathsf{Opt}$, which proves Lemma 3.2. Hence it suffices to prove Lemma 3.4. Note the difference from the induction used for set-cover: the thresholds used to define nets is non-uniform over the stages.

Proof: We induct on $j \in \{0, \cdots, T\}$. The base case is $j = T$, where we have a complete scenario-sequence $\mathbb{A}_T$: by feasibility of the optimum, $\phi^*_{\leq T}$ cuts $\widehat{A}_T \supseteq N_T$ from r in G. Thus the min-cut for N_T in $G \setminus \phi^*_{\leq T-1}(\mathbb{A}_{T-1})$ costs at most $c(\phi^*_T(\mathbb{A}_T)) = V_T(\mathbb{A}_T) \leq 5T \cdot V_T(\mathbb{A}_T)$.

Now assuming the inductive claim for $j + 1 \leq T$, we prove it for j. Let $H = G \setminus \phi^*_{\leq j-1}(\mathbb{A}_{j-1})$ be the residual graph after $j - 1$ stages, and let $H' = H \setminus \phi^*_j(\mathbb{A}_j)$ the residual graph after j stages. Let us divide up N_j into two parts, $N^1_j := \{v \in N_j \mid \mathsf{MinCut}_{H'}(v) > (2T - j - 1) \cdot \tau\}$ and $N^2_j = N_j \setminus N^1_j$, and bound the mincut of the two parts in H' separately.

Claim 3.5 $\mathsf{MinCut}_{H'}(N^2_j) \leq 4T \cdot c(\phi^*_j(\mathbb{A}_j))$.

Proof: Note that the set N^2_j consists of the points that have "high" mincut in the graph H after $j - 1$ stages, but have "low" mincut in the graph H' after j stages. For these we use a *Gomory-Hu tree*-based argument like that in [5]. Formally, let $t := (2T - j) \cdot \tau \leq 2T\tau$. Hence for every $u \in N^2_j$, we have:

$$\mathsf{MinCut}_H(u) > t \qquad \text{and} \qquad \mathsf{MinCut}_{H'}(u) \leq \left(1 - \tfrac{1}{2T}\right) t. \tag{3.3}$$

Consider the Gomory-Hu (cut-equivalent) tree $\mathcal{T}(H')$ on graph H', and root it at ρ. For each vertex $u \in N^2_j$, let $(X_u, \overline{X}_u)$ denote the minimum ρ-u cut in $\mathcal{T}(H')$, where $u \in X_u$ and $\rho \notin X_u$. Pick a subset $N' \subseteq N^2_j$ such that the union of their respective min-cuts in $\mathcal{T}(H')$ separate all of N^2_j from ρ and their corresponding sets X_u are disjoint—the set of cuts in tree $\mathcal{T}(H')$ closest to the root ρ gives such a collection. Define $F := \cup_{u \in N'} \partial_{H'}(X_u)$; this is a feasible cut in H' separating N^2_j from ρ.

Note that (3.3) implies that for all $u \in N^2_j$ (and hence for all $u \in N'$), we have

(i) $c(\partial_{H'}(X_u)) \leq (1 - \frac{1}{2T}) \cdot t$ since X_u is a minimum ρ-u cut in H', and
(ii) $c(\partial_H(X_u)) \geq t$ since it is a feasible ρ-u cut in H.

Thus $c(\partial_{H \setminus H'}(X_u)) = c(\partial_H(X_u)) - c(\partial_{H'}(X_u)) \geq \frac{1}{2T}\, t \geq \frac{1}{2T} \cdot c(\partial_{H'}(X_u))$. So

$$c(\partial_{H'}(X_u)) \leq 2T \cdot c(\partial_{H \setminus H'}(X_u)) \quad \text{for all } u \in N', \tag{3.4}$$

Consequently,

$$c(F) \leq \textstyle\sum_{u \in N'} c(\partial_{H'}(X_u)) \leq 2T \cdot \sum_{u \in N'} c(\partial_{H \setminus H'}(X_u)) \leq 4T \cdot c(H \setminus H') = 4T \cdot c(\phi^*_j(\mathbb{A}_j)).$$

The first inequality follows from subadditivity, the second from (3.4), the third uses disjointness of $\{X_u\}_{u \in N'}$, and the equality follows from $H \setminus H' = \phi^*_j(\mathbb{A}_j)$. Thus $\mathsf{MinCut}_{H'}(N^2_j) \leq 4T \cdot c(\phi^*_j(\mathbb{A}_j))$. ◀

Now to bound the cost of separating N_j^1 from ρ. Recall the quantity $W_j(\mathbb{A}_j)$ from (2.2),

$$W_j(\mathbb{A}_j) := \max_{|B_{j+1}|=k_{j+1}} V_{j+1}(A_1, \ldots, A_j, B_{j+1}).$$

and that $V_j(\mathbb{A}_j) = \phi_j^*(\mathbb{A}_j) + r \cdot W_j(\mathbb{A}_j)$.

Claim 3.6 $\mathsf{MinCut}_{H'}(N_j^1) \leq 5r\,T \cdot W_j(\mathbb{A}_j)$.

Proof: The definition of N_j^1 implies that for each $u \in N_j^1$ we have:

$$\mathsf{MinCut}_{H'}(u) \;>\; (2T - j - 1)\tau \;\geq\; T\,\tau \;\geq\; T \cdot \beta\,\frac{\mathsf{Opt}}{\lambda_{j+1}\,k_{j+1}} \;\geq\; \beta T\,\frac{W_j(\mathbb{A}_j)}{k_{j+1}}, \tag{3.5}$$

where the last inequality is by Claim 2.8. Furthermore, for any k_{j+1}-subset $L \subseteq N_j^1 \subseteq A_j$ we have:

$$\mathsf{MinCut}_{H'}(L) \;\leq\; 5T \cdot V_{j+1}(A_1, \cdots, A_j, L) \;\leq\; 5T \cdot W_j(\mathbb{A}_j). \tag{3.6}$$

The first inequality is by applying the induction hypothesis to $(A_1, \cdots, A_j, L)$; induction can be applied since L is a "net" for this partial scenario sequence (recall $L \subseteq N_j^1$ and the definition of N_j^1). The second inequality is by definition of $W_j(\mathbb{A}_j)$.

Now we apply Lemma 3.3 on graph H' with terminals $X = N_j^1$, bound $B = 5T \cdot W_j(\mathbb{A}_j)$, and $k = k_{j+1}$. Since $\beta = 50$, equations (3.5)-(3.6) imply that the conditions in Lemma 3.3 are satisfied, and we get $\mathsf{MinCut}_{H'}(N_j^1) \leq 5r\,T \cdot W_j(\mathbb{A}_j)$ to prove Claim 3.6. ◀

Finally,

$$\begin{aligned}\mathsf{MinCut}_H(N_j) \;&\leq\; \mathsf{MinCut}_{H'}(N_j^1) + \mathsf{MinCut}_{H'}(N_j^2) + c(\phi_j^*(\mathbb{A}_j)) \\ &\leq\; 5rT \cdot W_j(\mathbb{A}_j) + 4T\,c(\phi_j^*(\mathbb{A}_j)) + c(\phi_j^*(\mathbb{A}_j)) \;\leq\; 5T \cdot V_j(\mathbb{A}_j).\end{aligned}$$

The first inequality uses subadditivity of the cut function, the second uses Claims 3.5 and 3.6, and the third uses $T \geq 1$ and definition of $W_j(\mathbb{A}_j)$. This completes the proof of the inductive step, and hence of Lemma 3.4. ■

4 Multistage Robust Steiner Tree

We now turn to the multistage robust Steiner tree problem.

Theorem 4.1 *There is an $O(\min\{T, \log n, \log \lambda_{max}\})$-approximation algorithm for T-stage k-robust Steiner tree.*

Again, we prove an $O(T)$-approximation ratio where T is the number of stages; Theorem 4.1 then follows from the scaling arguments in Appendix A. This is the first of the problems we consider where the net creation is not "parallel": whereas in the two previous problems, we set the net to be all elements that were heavy in some formal sense, here we will pick as our net a set of elements that are mutually far from each other.

The input consists of an undirected edge-weighted graph $G = (U, E)$. For any edge-weighted graph $G' = (U', E')$ and $u, v \in U'$, let $d_{G'}(u, v)$ denote the shortest-path distance between u and v in G'; if no graph is specified then the distance is relative to the original G. Given a graph $G' = (U', E')$ and a set of edges $E'' \subseteq E'$, the graph G/E'' is defined by resetting the edges in E'' to have zero length. For any graph G' and subset X of vertices, $\mathsf{MinSt}_{G'}(X)$ is the minimum cost of a Steiner tree

on X. Recall the definition of a scenario sequence $\mathbb{A} = (A_0, A_1, \ldots, A_T)$, partial scenario sequence $\mathbb{A}_j = (A_0, A_1, \ldots, A_j)$ and notation from Section 2. The optimal strategy is $\Phi^* = \{\phi_j^*\}_{j=0}^T$, where $\phi_j^*(\mathbb{A}_j)$ maps to a set of edges in G to be chosen in stage j. The feasibility constraint is that $\phi_{\le T}^*(\mathbb{A}_T)$ connect the vertices in $\cap_{i \le T} A_i$ to each other. The optimal cost is $\mathsf{Opt} = \mathsf{RobCov}(\Phi^*)$.

Define $\tau := \beta \cdot \max_{j \in [T]} \frac{\mathsf{Opt}}{\lambda_j k_j}$ for $\beta = 10$; and $j^* = \operatorname{argmin}_{j \in [T]}(\lambda_j k_j)$. Let N be a maximal subset of U such that $d(u, v) > 4T \cdot \tau$ for all $u \neq v$, $u, v \in N$. The algorithm is:

> On day 0, buy edges $\phi_0 := \mathsf{MST}_G(N)$ i.e. a 2-approximate Steiner tree.
>
> On day j^*, for each $u \in A_{j^*}$, buy a shortest path from u to N in the residual graph G/ϕ_0.
>
> On all other days, do nothing.

Again, the cost incurred on day j^* is not high: by the maximality of N, the distance from any $u \in A_{j^*}$ to N is at most $4T\tau$, and hence the total effective cost incurred is $4T\tau \cdot |A_{j^*}| \cdot \lambda_{j^*} = 4T\beta \cdot \mathsf{Opt} = O(T)\,\mathsf{Opt}$. Thus it is enough to prove the following:

Lemma 4.2 *The MST on the set N costs at most $O(T) \cdot \mathsf{Opt}$.*

Let us define some useful notation, and review some concepts from previous sections. For any $v \in U$, distance $\delta \in \mathbb{R}_+$ and graph G', let the ball $\mathsf{B}_{G'}(v, \delta)$ denote the set of vertices that are at most distance δ from v. Given some value $t \ge 0$, a set $N \subseteq U$ of vertices is called a *t-ball-packing* in graph G' if the balls $\{\mathsf{B}_{G'}(u, t)\}_{u \in N}$ are disjoint. Finally, recall $V_j(\mathbb{A}_j)$ as defined in Definition 2.5; here we set $r = 1$. Given $\mathbb{A}$, recall that $\widehat{A}_j = \cap_{i \le j} A_i$.

The main structure lemma that will prove Lemma 4.2 is the following:

Lemma 4.3 *For any $j \in [T]$ and partial scenario-sequence $\mathbb{A}_j$, if the residual graph $H := G/\phi_{\le j-1}^*(\mathbb{A}_{j-1})$ and if $N_j \subseteq \widehat{A}_j$ is* **any** *$(2T - j)\tau$-ball-packing in H (i.e., the "net" elements) then the cost of the minimum Steiner tree $\mathsf{MinSt}_H(N_j) \le 5T \cdot V_j(\mathbb{A}_j)$.*

Before we prove Lemma 4.3, note that setting $j = 0$, the lemma says that the weight of the minimum Steiner tree in G connecting up elements of N (which is a $2T\tau$-ball-packing) costs at most $5T\,V_0 \le 5T\,\mathsf{Opt}$, which proves Lemma 4.2. One difference from the proof strategy used earlier (set cover and min-cut) is that we cannot directly rely on structural properties from the two-stage problem [6], since this yields only a guarantee exponential in T. Instead we give a different self-contained proof of the inductive step to obtain an $O(T)$ approximation ratio; this also yields an alternate proof of constant approximation for 2-stage robust Steiner tree.

Now, back to the proof of the lemma.

Proof: Let $\tau_i = (2T - i)\tau$ for any $i \in [T]$. In the base case $j = T$, we given the complete scenario sequence $\mathbb{A}$. The feasibility of the optimal solution implies that the edges in $\phi_{\le T}^*(\mathbb{A})$ connect up $\widehat{A}_T$, and hence N_T. Consequently, the cost to connect up N_T in $G/\phi_{\le T-1}^*(\mathbb{A}_{t-1})$ is at most $c(\phi_T^*(\mathbb{A}_T)) = V_T(\mathbb{A}_T) \le 5T\,V_T(\mathbb{A}_T)$, since $T \ge 1$.

For the inductive step, assume that the lemma holds for $j + 1$, and let us prove it for j. Let $H = G/\phi_{\le j-1}^*(\mathbb{A}_{j-1})$, and $H' = H/\phi_j^*(\mathbb{A}_j) = G/\phi_{\le j}^*(\mathbb{A}_j)$. Note that by assumption, N_j is a τ_j-ball-packing in graph H, so the balls $\{\mathsf{B}_H(u, \tau_j) \mid u \in N_j\}$ are disjoint.

Define $N_j^1 := \{v \in N_j \mid \mathsf{B}_{H'}(v, \tau_{j+1}) \subseteq \mathsf{B}_H(v, \tau_j)\}$ be the set of points in N_j such that their τ_{j+1}-balls in H' are contained within their τ_j-balls in H. (Recall that $\tau_{j+1} < \tau_j$, and hence this can indeed happen.) Let $N_j^2 = N_j \setminus N_j^1$ be the remaining points in N_j. Let us prove some useful facts about these two sets.

Claim 4.4 $|N_j^2|\,\tau_{j+1} \le 2T \cdot c(\phi_j^*(\mathbb{A}_j))$.

Proof: For any $v \in N_j^2$, it holds that $\mathsf{B}_{H'}(v, \tau_{j+1}) \not\subseteq \mathsf{B}_H(v, \tau_j)$. So there is some vertex $w \in \mathsf{B}_{H'}(v, \tau_{j+1}) \setminus \mathsf{B}_H(v, \tau_j)$; this means that $d_H(v, w) > \tau_j$ but $d_{H'}(v, w) \le \tau_{j+1}$; the distance to w has shrunk by at least $\tau_j - \tau_{j+1} = \tau$. Since $H' = H/\phi_j^*(\mathbb{A}_j)$, this implies that the edges bought by $\phi_j^*(\mathbb{A}_j)$ just within this ball is large. In particular, if $E(\mathsf{B}_H(v, \tau_j))$ denotes the edges induced on $\mathsf{B}_H(v, \tau_j)$ then:

$$c\left(\phi_j^*(\mathbb{A}_j) \bigcap E(\mathsf{B}_H(v, \tau_j))\right) \ge \tau_j - \tau_{j+1} = \tau, \quad \text{which holds for all } v \in N_j^2.$$

Since the edge-sets $E(\mathsf{B}_H(v, \tau_j))$ for $v \in N_j^2$ are disjoint, we can sum over all vertices in N_j^2 to get $c(\phi_j^*(\mathbb{A}_j)) \ge |N_j^2| \cdot \tau$. Finally, $\tau_{j+1} \le 2T \cdot \tau$, so $|N_j^2|\,\tau_{j+1} \le 2T \cdot c(\phi_j^*(\mathbb{A}_j))$ which proves the claim. ◀

Claim 4.5 *The set N_j^1 forms a τ_{j+1}-ball-packing in H'.*

Proof: Assume not. Then the τ_{j+1}-balls in H' around some two points $u, v \in N_j^1$ must intersect. But these balls are contained within the τ_j-balls around them in H, by the definition of N_j^1, so $\mathsf{B}_H(u, \tau_j) \cap \mathsf{B}_H(v, \tau_j) \neq \emptyset$. But this contradicts the fact that $u, v \in N_j$, and N_j was a τ_j-ball-packing in H. ◀

Recall the definition of $W_j(\mathbb{A}_j)$ as in (2.2); note that we set $r = 1$.

Claim 4.6 *For any τ_{j+1}-ball-packing $Z \subseteq N_j$ in H', its size $|Z| \le k_{j+1} - 1$. Moreover, $\mathsf{MinSt}_{H'}(Z) \le 5T \cdot W_j(\mathbb{A}_j)$.*

Proof: For a contradiction, suppose $|Z| \ge k_{j+1}$, and let $A_{j+1} \subseteq Z$ denote any k_{j+1}-set. Observe that $(A_1, \cdots, A_j, A_{j+1})$ is a valid partial scenario-sequence. Also, $A_{j+1} \subseteq Z \subseteq N_j \subseteq \widehat{A}_j$, and hence $A_{j+1} \subseteq \cap_{i \le j+1} A_i = \widehat{A}_{j+1}$. Furthermore, $A_{j+1} \subseteq Z$ is a τ_{j+1}-ball-packing in H'. Thus applying the induction hypothesis for $j+1$ on $(A_1, \cdots, A_j, A_{j+1})$, we get that $\mathsf{MinSt}_{H'}(A_{j+1}) \le 5T \cdot V_{j+1}(\mathbb{A}_{j+1})$. Now,

$$5T \cdot V_{j+1}(\mathbb{A}_{j+1}) \;\le\; 5T \cdot W_j(\mathbb{A}_j) \;\le\; 5T \cdot \frac{\mathsf{Opt}}{\lambda_{j+1}} \;\le\; \frac{5T}{\beta} \cdot k_{j+1}\,\tau \;\le\; \frac{T\,k_{j+1}\,\tau}{2}.$$

The first inequality is by (2.2) which defines $W_j(\mathbb{A}_j)$, the second by Claim 2.8, the third by the definition of τ, and the last inequality uses $\beta = 10$. On the other hand, A_{j+1} is a τ_{j+1}-ball-packing in H' and so $\mathsf{MinSt}_{H'}(A_{j+1}) > (|A_{j+1}| - 1)\tau_{j+1} = (k_{j+1} - 1) \cdot \tau_{j+1} \ge k_{j+1} T\,\tau/2$, since we may assume $k_{j+1} \ge 2$ (otherwise $k_T \le 1$ and the optimal value is 0). This contradicts the above bound. Thus we must have $|Z| < k_{j+1}$.

Now augment Z with any $k_{j+1} - |Z|$ elements to obtain A_{j+1}, and apply the inductive hypothesis again on the scenario-sequence $(A_1, \cdots, A_j, A_{j+1})$ and the τ_{j+1}-ball-packing in H', $Z \subseteq \widehat{A}_{j+1}$ to obtain

$$\mathsf{MinSt}_{H'}(Z) \;\le\; 5T \cdot V_{j+1}(\mathbb{A}_{j+1}) \;\le\; 5T \cdot W_j(\mathbb{A}_j). \tag{4.7}$$

This completes the proof of the claim. ◀

Construct a maximal τ_{j+1}-ball-packing $Z \subseteq N_j$ in the graph H' as follows: add all of N_j^1 to Z, and greedily add vertices from N_j^2 to Z until no more can be added without violating the ball-packing condition. Now, to prove the inductive step, we need to show how to connect up the set N_j cheaply in the graph H. We first bound the cost of this Steiner tree in H'. Claim 4.6 says we can connect up Z in the graph H' at cost $5T\,W_j(\mathbb{A}_j)$. Since Z is a maximal τ_{j+1}-ball-packing inside N_j, we know that each element in $N_j \setminus Z \subseteq N_j^2$ is at distance at most $2\tau_{j+1}$ from some element of Z. There are

only $2Tc(\phi_j^*(\mathbb{A}_j))/\tau_{j+1}$ elements in N_j^2 by Claim 4.4, so the total cost to connect these points to Z is at most their product $4Tc(\phi_j^*(\mathbb{A}_j))$, giving us that

$$\mathsf{MinSt}_{H'}(N_j) \leq 5T \cdot W_j(\mathbb{A}_j) + 4T \cdot c(\phi_j^*(\mathbb{A}_j))$$

Finally, since the length of the Steiner-tree(N_j) in H can only be greater by $\phi_j^*(\mathbb{A}_j)$, we get that

$$\mathsf{MinSt}_{H}(N_j) \leq 5T \cdot W_j(\mathbb{A}_j) + 5T \cdot c(\phi_j^*(\mathbb{A}_j)) = 5T \cdot V_j(\mathbb{A}_j),$$

which completes the proof of the lemma. ∎

5 Multistage Robust Steiner Forest

Here we consider the multistage robust version of the Steiner Forest problem:

Theorem 5.1 *There is an $O(\min\{T, \log n, \log \lambda_{max}\})$-approximation algorithm for T-stage k-robust Steiner forest.*

Again, we prove an $O(T)$-approximation ratio where T is the number of stages; Theorem 5.1 then follows from the scaling arguments in Appendix A. Here, both the algorithm and the proof will be slightly more involved than the previous problems. Recall that the input for Steiner Forest consists of an undirected edge-weighted graph G, and pairs $\{s_i, t_i\}_{i \in P}$. For any graph H and subset $P' \subseteq P$ of pairs, we let $\mathsf{MinSF}_H(P')$ denote the minimum cost of a Steiner forest connecting pairs in P'; again if the graph is not specified it is relative to the input graph G.

For $\Delta \geq 0$, we define a subset $N \subseteq P$ of pairs to be a Δ-*SFnet* in graph G' if

- $d_{G'}(s_i, t_i) > \Delta$ for all $i \in N$, and
- there exist $z_i \in \{s_i, t_i\}$ for all $i \in N$ such that the balls $\{\mathsf{B}_{G'}(z_i, \Delta)\}_{i \in N}$ are disjoint.

The following simple property immediately follows from the definition:

Lemma 5.2 *If $N \subseteq P$ is a Δ-SFnet in graph G' then $\mathsf{MinSF}_{G'}(N) \geq |N| \cdot \Delta$.*

Proof: Consider the dual of the Steiner forest LP relaxation, which is a packing problem. Each ball of radius at most Δ around any vertex in $\{z_i\}_{i \in N}$ is a feasible variable in the dual, since $d_{G'}(s_i, t_i) > \Delta$ for all $i \in N$. Now since $\{\mathsf{B}_{G'}(z_i, \Delta)\}_{i \in N}$ are disjoint, there is a feasible dual solution of value $|N| \cdot \Delta$. Hence the optimal Steiner forest costs at least as much. ∎

However, it seems quite non-trivial to even compute a maximal Δ-SFnet. This is unlike the previous algorithms (set cover, min-cut, Steiner tree) where computing this net was straightforward. Instead, we will run a specialized procedure to compute a near-maximal net which suffices to give an algorithm for multistage Steiner forest.

Now to describe the algorithm. Define $\beta := 10$, $\tau := \beta \cdot \max_{j \in [T]} \frac{\mathsf{Opt}}{\lambda_j k_j}$ and $j^* = \operatorname{argmin}_{j \in [T]}(\lambda_j k_j)$. We now run Algorithm 5.1 below to find a $\gamma = 2T\tau$-SFnet $N \subseteq P$, as well as a set of edges $E_{\mathsf{alg}} \subseteq E$. Then,

> On day 0, buy the edges in $\phi_0 := E_{\mathsf{alg}}$.
>
> On day j^*, for each $(s, t) \in A_{j^*}$, buy a shortest (s, t)-path in the residual graph G/ϕ_0.
>
> On all other days, do nothing.

Algorithm 5.1 Algorithm for near-maximal γ-SFnet for Steiner forest

1: **let** $S_r, S_g, S_o, S_b, S_f, W \leftarrow \emptyset$.
2: **while** there exists a pair $i \in P$ with $d_{G/(S_r \cup S_f)}(s_i, t_i) > 4\gamma$ **do**
3: **let** $S_r \leftarrow S_r \cup \{i\}$
4: **if** $d_G(s_i, w) < 2\gamma$ for some $w \in W$ **then** $S_f \leftarrow S_f \cup \{(s_i, w)\}$ **else** $W \leftarrow W \cup \{s_i\}$
5: **if** $d_G(t_i, w') < 2\gamma$ for some $w' \in W$ **then** $S_f \leftarrow S_f \cup \{(t_i, w')\}$ **else** $W \leftarrow W \cup \{t_i\}$
6: **let** $\delta_i \in \{0, 1, 2\}$ be the increase in $|W|$ due to steps (4) and (5).
7: **if** $(\delta_i = 0)$ **then** $S_b \leftarrow S_b \cup \{i\}$; **if** $(\delta_i = 1)$ **then** $S_o \leftarrow S_o \cup \{i\}$; **if** $(\delta_i = 2)$ **then** $S_g \leftarrow S_g \cup \{i\}$.
8: **let** $E_{\mathsf{alg}} :=$ 2-approximate Steiner forest on pairs S_r, along with shortest-paths connecting every pair in S_f.
9: **output** the γ-SFnet $N := S_g \bigcup S_o$, and edge set E_{alg}.

Set $\gamma := 2T\tau$. In Algorithm 5.1, $G/(S_r \cup S_f)$ denotes the graph obtained from G by identifying all pairs in S_r and S_f. We note that this algorithm is essentially same as the one for 2-stage robust Steiner forest in [6]; however we need to rephrase slightly in order to fit it into our context.

To bound the cost incurred on day j^*, we note that the algorithm adds pairs in P whose distance in the residual graph G/E_{alg} is more than $4\gamma = 8T\tau$. Hence the cost of connecting any pair in $P \setminus N$ is at most $8T\tau$, and consequently the total effective cost incurred on day j^* is $8T\tau \cdot |A_{j^*}| \cdot \lambda_{j^*} = 8T\beta \cdot \mathsf{Opt} = O(T)\,\mathsf{Opt}$. Thus it is enough to prove the following lemma, whose proof we present in the next section:

Lemma 5.3 *The cost of the edges E_{alg} is at most $144T\,\mathsf{Opt}$.*

The proof bounding the cost of edges in E_{alg} is also more complicated that in previous sections; here's the roadmap. First, we use the (by now familiar) reverse inductive proof to bound the cost of edges in an optimal Steiner Forest connecting the pairs in the net N. However, our set of edges E_{alg} is not just an approximate Steiner Forest for the pairs in N, but also connects other nodes, and we have to also bound the cost of the extra edges our algorithm buys.

For the rest of the section, we use $\gamma := 2T\tau$. Recall that $\widehat{A}_j = \cap_{i \le j} A_j$. Also recall $V_j(\mathbb{A}_j)$ from Definition 2.5; here we use $r = 2$. Recall that a set $S \subseteq U$ in a graph $G' = (U, E')$ is called a *t-ball-packing* if the balls $\{\mathsf{B}_{G'}(x, t)\}_{t \in S}$ are disjoint. Hence, the second condition in the definition of a set N being a Δ-SFnet is that for each $\{s_i, t_i\} \in N$, there exists $z_i \in \{s_i, t_i\}$ such that the set of these z_i's is a Δ-ball-packing. If N is a SFnet, let $\zeta(N)$ denote this ball-packing that witnesses it.

Lemma 5.4 *For any $j \in [T]$ and partial scenario-sequence $\mathbb{A}_j$, let the residual graph in stage j of the optimal algorithm be $H := G/\phi^*_{\le j-1}(\mathbb{A}_{j-1})$, and let $N_j \subseteq \widehat{A}_j$ be* **any** *$(2T - j)\tau$-SFnet in H (i.e., the "net" elements). Then the cost of the optimal Steiner forest on N_j is $\mathsf{MinSF}_H(N_j) \le 9T \cdot V_j(\mathbb{A}_j)$.*

Proof: Let $\tau_i = (2T - i)\tau$ for any $i \in [T]$. In the base case $j = T$, we are given the complete scenario sequence $\mathbb{A}$. The feasibility of the optimal solution implies that the edges in $\phi^*_{\le T}(\mathbb{A})$ connect the pairs $\widehat{A}_T$, and hence N_T. Consequently, the cost to connect up N_T in $G/\phi^*_{\le T-1}(\mathbb{A}_{t-1})$ is at most $c(\phi^*_T(\mathbb{A}_T)) = V_T(\mathbb{A}_T) \le 9T \cdot V_T(\mathbb{A}_T)$.

Let $\mathcal{T}_j := \{s_i, t_i \mid i \in N_j\}$ be all the terminals in N_j. Since N_j is a τ_j-SFnet in graph H, $\zeta(N_j)$ is a τ_j-ball-packing in H. Let graph $H' := H/\phi^*_j(\mathbb{A}_j)$. Similar to the proof for Steiner tree, define $N^1_j := \{i \in N_j \mid \mathsf{B}_{H'}(z_i, \tau_{j+1}) \subseteq \mathsf{B}_H(z_i, \tau_j)\}$, and let $N^2_j = N_j \setminus N^1_j$ be the rest of the pairs. The fact that N^1_j is a τ_{j+1}-SFnet in H' follows from the definition of N^1_j and Claim 4.5. We bound the costs

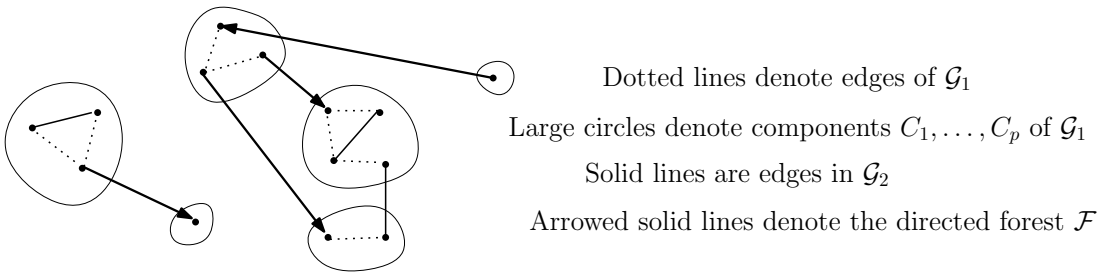

Figure 5.1: The auxiliary graphs $\mathcal{G}_1$ and $\mathcal{G}_2$.

$\mathsf{MinSF}_{H'}(N_j^1)$ and $\mathsf{MinSF}_{H'}(N_j^2)$ separately by invoking the inductive hypothesis twice. Let $W_j(\mathbb{A}_j)$ be defined as in (2.2); note that $V_j(\mathbb{A}_j) = c(\phi_j^*(\mathbb{A}_j) + 2\,W_j(\mathbb{A}_j)$, since $r = 2$ in this case.

The next two claims have proofs almost identical to the Steiner tree Claim 4.6, and Claim 4.4 respectively.

Claim 5.5 *The cardinality of N_j^1 is $|N_j^1| < k_{j+1}$, and $\mathsf{MinSF}_{H'}(N_j^1) \le 9T \cdot W_j(\mathbb{A}_j)$.*

Claim 5.6 *We have $|N_j^2|\,\tau_{j+1} \le 2T \cdot c(\phi_j^*(\mathbb{A}_j))$.*

Recall we need to show that the cost to connect up pairs in N is small; we now have that connecting up N_j^1 is not expensive, and have a bound on the cardinality of N_j^2—it remains to show that this can be used to bound its cost. This we do in the following discussion culminating in Claim 5.9.

Let $\mathcal{T}_j^2 := \{s_i, t_i \mid i \in N_j^2\}$ be the terminals corresponding to pairs in N_j^2. Consider an auxiliary graph $\mathcal{G}_1$ on these terminals, where edges connect nodes at distance at most $2\tau_{j+1}$—i.e., $V(\mathcal{G}_1) = \mathcal{T}_j^2$, and $E(\mathcal{G}_1) = \{(a,b) \mid a, b \in \mathcal{T}_j^2,\, d_{H'}(a,b) \le 2\tau_{j+1}\}$. Let $C_1, \ldots, C_p$ denote the connected components in this graph $\mathcal{G}_1$, so $\mathcal{T}_j^2 = \cup_{\ell=1}^{p} C_\ell$. See also Figure 5.1. Let $\mathcal{M}$ denote the minimum length forest in the graph H' having the same connected components as $\mathcal{G}_1$.

Claim 5.7 *The cost of the forest $\mathcal{M}$ (in graph H') is at most $8T \cdot c(\phi_j^*(\mathbb{A}_j))$.*

Proof: Two nodes in $\mathcal{G}_1$ are connected only if their distance in H' is at most $2\tau_{j+1}$; moreover, the number of vertices in $\mathcal{G}_1$ is $|\mathcal{T}_j^2| = 2|N_j^2|$. Hence, the forest $\mathcal{M}$ costs at most $4|N_j^2|\tau_{j+1}$, which is at most $8T\ c(\phi_j^*(\mathbb{A}_j))$ using Claim 5.6. ◀

Define $\alpha : \mathcal{T}_j^2 \to \{1, \ldots, p\}$ mapping each terminal in $\mathcal{T}_j^2$ to the index for its component in $\mathcal{G}_1$: i.e., $\alpha(u) = \ell$ if $u \in C_\ell \subseteq \mathcal{G}_1$. Define another auxiliary multigraph $\mathcal{G}_2$ on vertices $\{1, 2, \ldots, p\}$ and edges $E(\mathcal{G}_2) := \{(\alpha(s_i), \alpha(t_i)) \mid i \in N_j^2\}$. (See Figure 5.1.) Observe that there is a one-to-one correspondence between edges in $\mathcal{G}_2$ and pairs in N_j^2. Let $\mathcal{F}$ denote the edges of any maximal forest in $\mathcal{G}_2$ ($\mathcal{F}$ cannot contain any self-loops); we also use $\mathcal{F}$ to denote the corresponding pairs from N_j^2.

Claim 5.8 $\mathsf{MinSF}_{H'}(\mathcal{F}) \le 9T \cdot W_j(\mathbb{A}_j)$.

Proof: Orient the edges in the forest $\mathcal{F}$ so that each vertex $\ell \in V(\mathcal{G}_2)$ has at most one incoming edge; call this map $\pi : \mathcal{F} \to V(\mathcal{G}_2)$, and note that for each $\ell \in \mathcal{F}$, $\pi^{-1}(\ell) \le 1$. For each $i \in \mathcal{F}$, set $z_i' = s_i$ if the edge is oriented towards s_i, and t_i otherwise. Clearly, the vertices $\{z_i'\}_{i \in \mathcal{F}}$ lie in distinct components of $\mathcal{G}_1$. Hence $\{z_i'\}_{i \in \mathcal{F}}$ is an independent set in $\mathcal{G}_1$, and hence is a τ_{j+1}-ball-packing in H'. Moreover, since $\{s_i, t_i\}$ lie in distinct components of $\mathcal{G}_1$, $d_{H'}(s_i, t_i) > 2\tau_{j+1}$; hence $\mathcal{F}$ is also a τ_{j+1}-SFnet in graph H'. Now, as in the proof of Claim 4.6, it follows by induction on $\mathcal{F}$ that $\mathsf{MinSF}_{H'}(\mathcal{F}) \le 9T \cdot W_j(\mathbb{A}_j)$. ◀

Claim 5.9 $\mathsf{MinSF}_{H'}(N_j^2) \le \mathsf{MinSF}_{H'}(\mathcal{F}) + c(\mathcal{M}) \le 9T \cdot W_j(\mathbb{A}_j) + 8T \cdot c(\phi_j^*(\mathbb{A}_j))$.

Proof: We will show that for every $i \in N_j^2$, the subgraph $\mathcal{M} \bigcup \mathsf{MinSF}_{H'}(\mathcal{F})$ connects s_i to t_i in graph H', after which the claim follows directly from Claims 5.7 and 5.8.

Recall that $\mathcal{M}$ has the same connected components as graph $\mathcal{G}_1$. For any $i \in N_j^2$, let $a = \alpha(s_i)$ and $b = \alpha(t_i)$ be the indices of the components containing s_i and t_i respectively. If $a = b$, then $\mathcal{M}$ connects $\{s_i, t_i\}$, so assume $a \neq b$. Since $\mathcal{F}$ is a maximal forest in $\mathcal{G}_2$, it contains some path $e_1, \ldots, e_q$ from $\alpha(a)$ to $\alpha(b)$. But now $\mathsf{MinSF}_{H'}(\mathcal{F}) \cup \mathcal{M}$ connects s_i to t_i. ◀

Finally using sub-additivity of Steiner Forest, and Claims 5.5 and 5.9,

$$\mathsf{MinSF}_{H'}(N_j) \le \mathsf{MinSF}_{H'}(N_j^1) + \mathsf{MinSF}_{H'}(N_j^2) \le 2 \cdot 9T \cdot W_j(\mathbb{A}_j) + 8T \cdot c(\phi_j^*(\mathbb{A}_j)).$$

Thus $\mathsf{MinSF}_H(N_j) \le c(\phi_j^*(\mathbb{A}_j)) + \mathsf{MinSF}_{H'}(N_j) \le 9T \cdot [c(\phi_j^*(\mathbb{A}_j)) + 2 \cdot W_j(\mathbb{A}_j)] = 9T \cdot V_j(\mathbb{A}_j)$. ■

As a consequence of the case $j = 0$ of Lemma 5.4 from the previous section, we know that the cost of the optimal Steiner forest on any $\tau_0 = 2T\tau$-SFnet N_0 in the original graph G is at most $9T\,V_0 \le 9T\,\mathsf{Opt}$. And indeed, the set $N = S_g \cup S_o$ is a $2T\tau$-SFnet in G, since $\{\mathsf{B}_G(w, 2T\tau) \mid w \in W\}$ are disjoint and $W \cap \{s_i, t_i\} \neq \emptyset$ for all $i \in S_g \cup S_o$. However, the set E_{alg} is not just the Steiner forest on $S_g \cup S_o$, it is actually a Steiner forest on $S_r = S_g \cup S_o \cup S_b$, along with shortest paths between every "fake" pair in S_f. This is what we bound in the proof below.

Proof of Lemma 5.3: Observe that $c(E_{\mathsf{alg}}) \le 2 \cdot \mathsf{MinSF}(S_r) + 2\gamma \cdot |S_f|$, since the distance between each pair in S_f is at most 2γ. (Remember, $\gamma = 2T\tau$.) Recall that $N = S_g \cup S_o$ and $S_r = N \cup S_b$; so $\mathsf{MinSF}(S_r) \le \mathsf{MinSF}(N) + \mathsf{MinSF}(S_b)$. Finally, the argument in the previous paragraph shows that $\mathsf{MinSF}_G(N) \le 9T\,V_0 \le 9T\,\mathsf{Opt}$. Hence

$$c(E_{\mathsf{alg}}) \le 2 \cdot 9T\,\mathsf{Opt} + 2 \cdot \mathsf{MinSF}(S_b) + 2\gamma \cdot |S_f|. \tag{5.8}$$

Moreover, S_b might be very far from a γ-SFnet, so we cannot just apply the same techniques to it.

Bounding $\mathsf{MinSF}(S_b)$. As in the proof of Lemma 5.4, define an auxiliary graph $\mathcal{G}_1$ with vertices $V(\mathcal{G}_1) = \{s_i, t_i \mid i \in S_b\}$, and edges $E(\mathcal{G}_1) = \{(a, b) \mid d_G(a, b) \le 2\gamma\}$ between any two terminals in $V(\mathcal{G}_1)$ that are at most 2γ apart. Let $C_1, \ldots, C_p$ be the connected components in this graph $\mathcal{G}_1$, and let $\mathcal{M}$ be the minimum cost spanning forest in G having the same connected components as $\mathcal{G}_1$. Since there are $2|S_b|$ vertices in $\mathcal{G}_1$ and edges correspond to pairs at most 2γ from each other, the cost $c(\mathcal{M}) \le 4|S_b|\,\gamma$.

Again, define map $\alpha : V(\mathcal{G}_1) \to \{1, \ldots, p\}$ where $\alpha(u) = \ell$ if $u \in C_\ell$. Define another auxiliary graph $\mathcal{G}_2$ on vertices $\{1, 2, \ldots, p\}$ with edges $E(\mathcal{G}_2) := \{(\alpha(s_i), \alpha(t_i)) \mid i \in S_b\}$. Let $\mathcal{F}$ denote the edges of any maximal forest in $\mathcal{G}_2$, and also to denote the corresponding pairs from S_b. By orienting $\mathcal{F}$ so that each vertex has indegree at most one, we obtain $z_i' \in \{s_i, t_i\}$ for all $i \in \mathcal{F}$ satisfying $|C \cap \{z_i'\}_{i \in \mathcal{F}}| \le 1$ for each component C of $\mathcal{G}_1$ (see Claim 5.8 for details). Thus $\{z_i\}_{i \in \mathcal{F}}$ is an independent set in $\mathcal{G}_1$, and so $\{\mathsf{B}_G(z_i, r) \mid i \in \mathcal{F}\}$ are disjoint. This implies that $\mathcal{F}$ is a γ-SFnet in G. Applying Lemma 5.4 on $\mathcal{F}$ (with $j = 0$) gives $\mathsf{MinSF}_G(\mathcal{F}) \le 9T \cdot \mathsf{Opt}$.

Finally, as in Claim 5.9 we obtain $\mathsf{MinSF}(S_b) \le \mathsf{MinSF}(\mathcal{F}) + c(\mathcal{M}) \le 9T \cdot \mathsf{Opt} + 4|S_b|\,\gamma$. Combining this with (5.8) we have:

$$c(E_{\mathsf{alg}}) \le 36T \cdot \mathsf{Opt} + 8|S_b|\,\gamma + 2|S_f|\,\gamma. \tag{5.9}$$

Bounding $|S_b|$ and $|S_f|$. We use the following property of Algorithm 1.

Claim 5.10 ([6]) *In any execution of Algorithm 1, $|S_b| \le |N|$ and $|S_f| \le 2|N|$.*

Proof: This follows directly from the analysis in [6]. Lemma 5.3 in that paper yields $|S_f| \leq |S_r|$. By the definition of various sets S_g, S_b, S_o, S_f and $N = S_g \cup S_o$, we have that $|S_f| \geq 2|S_b|$ and $|S_r| = |N| + |S_b|$. Thus we have $2|S_b| \leq |S_f| \leq |S_r| \leq |N| + |S_b|$, implying $|S_b| \leq |N|$. Finally, $|S_f| \leq |S_r| = |N| + |S_b| \leq 2|N|$. ◀

Combining this with (5.9), it just remains to bound $|N|\,\gamma$. Recall that N is a γ-SFnet in G; so $\gamma\,|N| \leq \mathsf{MinSF}_G(N)$. On the other hand, we already argued that $\mathsf{MinSF}_G(N) \leq 9T\,\mathsf{Opt}$. Putting this together with (5.9), we get

$$c(E_{\mathsf{alg}}) \leq 36T \cdot \mathsf{Opt} + 8|S_b|\,\gamma + 2|S_f|\,\gamma \leq 36T \cdot \mathsf{Opt} + 12|N|\,\gamma \leq 144T \cdot \mathsf{Opt}.$$

This completes the proof of Lemma 5.3. ■

Acknowledgments. We thank F. B. Shepherd, A. Vetta, and M. Singh for their generous hospitality during the initial stages of this work.

References

[1] A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski. Adjustable robust solutions of uncertain linear programs. *Mathematical Programming*, 99(2):351–376, 2004.

[2] D. Bertsimas, D.B. Brown, and C. Caramanis. Theory and applications of robust optimization. *SIAM Review*, 53(3):464–501, 2011.

[3] K. Dhamdhere, V. Goyal, R. Ravi, and M. Singh. How to pay, come what may: Approximation algorithms for demand-robust covering problems. In *FOCS*, pages 367–378, 2005.

[4] U. Feige, K. Jain, M. Mahdian, and V.S. Mirrokni. Robust combinatorial optimization with exponential scenarios. In *IPCO*, pages 439–453, 2007.

[5] D. Golovin, V. Goyal, and R. Ravi. Pay today for a rainy day: improved approximation algorithms for demand-robust min-cut and shortest path problems. In *STACS*, pages 206–217. 2006.

[6] A. Gupta, V. Nagarajan, and R. Ravi. Thresholded Covering Algorithms for Robust and Max-Min Optimization (Full version: CoRR abs/0912.1045). In *ICALP*, pages 262–274, 2010.

[7] A. Gupta, M. Pál, R. Ravi, and A. Sinha. What About Wednesday? Approximation Algorithms for Multistage Stochastic Optimization. In *APPROX-RANDOM*, pages 86–98, 2005.

[8] R. Khandekar, G. Kortsarz, V.S. Mirrokni, and M.R. Salavatipour. Two-stage robust network design with exponential scenarios. In *ESA*, pages 589–600, 2008.

[9] C. Swamy and D. B. Shmoys. Sampling-based approximation algorithms for multistage stochastic optimization. *SIAM J. Comput.*, 41(4):975–1004, 2012.

A Preprocessing to Bound Number of Stages T

We show using a scaling argument how to ensure $T \leq O(\min\{\log \lambda_{max}, \log n\})$ for the min-cut, Steiner tree and Steiner forest problems; here n is the number of vertices in the input graph). Recall that the approximation for multistage robust set cover does not depend on T, hence such a preprocessing step is not required for set cover.

It is easy to ensure that $\lambda_{max} = \lambda_T \geq 2^T$ losing a constant factor in the objective: we simply find a maximal subsequence of stages i where the inflation λ_i in each stage in the subsequence increases by at least a factor of two from the previous stage in the subsequence, and perform actions only on these stages (or in the very final stage T). This means $T \leq \log_2 \lambda_{max}$.

We now ensure that $T \leq O(\log n)$. Guess the edge f of maximum cost (unscaled) that is ever bought by the optimal strategy on any scenario-sequence; there are at most n^2 choices for f and the algorithm will enumerate over these. It follows that $\mathsf{Opt} \geq c_f$ since f must be bought on some scenario-sequence, and all inflation factors are ≥ 1. Let $E_{high} := \{e \in E : c_e > c_f\}$; by the choice of f, no edge in E_{high} is ever used by Opt for any scenario-sequence. So the optimal value does not change even if we remove edges E_{high} from the instance—in Steiner tree/forest this means deleting edges E_{high} from the graph, and in min-cut this means contracting edges E_{high}. This yields an instance with maximum edge-cost $c_{max} \leq c_f$.

Let $E_{low} := \{e \in E : c_e < c_f/n^2\}$; the total cost of edges in E_{low} is at most $|E_{low}| \cdot \frac{c_f}{n^2} \leq c_f \leq \mathsf{Opt}$ by the choice of f. So we can assume that all edges in E_{low} are always bought in stage 0, at the loss of an additive Opt term in the objective. This ensures that the minimum edge-cost $c_{min} \geq c_f/n^2$. Combined with the above step, we have an instance with $\frac{c_{max}}{c_{min}} \leq n^2$.

Observe that any stage i with $\lambda_i > n^2 \cdot \frac{c_{max}}{c_{min}}$ is completely inactive in any optimal strategy: if not then the objective of the resulting strategy is greater than $n^2 \cdot c_{max}$, whereas the trivial strategy of buying all elements in stage 0 costs at most $n^2 \cdot c_{max}$. Hence, without loss of generality, we have $\lambda_{max} \leq n^2 \cdot \frac{c_{max}}{c_{min}}$, which combined with the above gives $\lambda_{max} \leq n^4$. Thus $T \leq \log_2 \lambda_{max} \leq O(\log n)$.

B Some Useful Examples

B.1 Non-Optimality of Two-Stage Strategies for Set Cover

We give an instance of multistage set-cover where any optimal solution has to buy sets on all days. This shows that we really need to consider near-optimal strategies to prove our structure result about thrifty strategies.

Let T denote the time horizon. The scenario bounds are $k_i = T + 1 - i$ for each $0 \leq i \leq T$. The inflation factors are $\lambda_i = (1+\epsilon)^i$ for each $i \in [T]$, where $\epsilon > 0$ is chosen such that $\lambda_T < 2$. The elements in the set-system are $U := \{1, \cdots, T, T+1\}$. The sets and their costs are as follows:

- For each $i \in \{1, \cdots, T-1\}$, define $\mathcal{S}_i$ as the collection consisting of all $k_i = T+1-i$ subsets of $\{i, \cdots, T+1\}$ *except* $\{i+1, \cdots, T+1\}$; each of the sets in $\mathcal{S}_i$ has cost λ_T/λ_i. Note that each set in $\mathcal{S}_i$ contains i.
- There are two singleton sets $\{T\}$ and $\{T+1\}$ of cost one each; let $\mathcal{S}_T := \{\{T\}, \{T+1\}\}$.

The number of sets is at most T^2, and each costs at least one.

Consider the strategy σ that does nothing on day 0 and for each day $i \in \{1, 2, \ldots, T\}$ does:

- If the scenario $A_i \in \binom{U}{k_i}$ in day i equals one of the sets in $\mathcal{S}_i$ then buy set A_i on day i. Note that this set A_i in the set-system has cost λ_T/λ_i.
- For any other scenario $A_i \in \binom{U}{k_i}$, buy nothing on day i.

It can be checked directly that this is indeed a feasible solution. Moreover, the effective cost under every scenario-sequence is exactly $\lambda_T < 2$. Note that for every scenario-sequence, σ buys sets on exactly one day; however the days corresponding to different scenarios are different. In fact, for each $i \in [T]$ there is a scenario-sequence (namely $A_j = \{j+1, \cdots, T+1\}$ for $j < i$ and $A_i \in \mathcal{S}_i$) where σ buys sets on day i. Thus strategy σ buys sets on all days.

We now claim that σ is the unique optimal solution to this instance. For any other feasible strategy σ', consider a scenario-sequence $(A_1, \cdots, A_T)$ where σ' behaves differently from σ. Let $i \in \{0, \cdots, T\}$ denote the earliest day when σ' differs from σ under this scenario-sequence. There are two possibilities:

- $A_j = \{j+1, \cdots, T+1\}$ for each $j < i$ and $A_i \in \mathcal{S}_i$, i.e. σ buys nothing before day i and buys set A_i on day i. σ' buys nothing before day i and covers $A'_i \subsetneq A_i$ on day i. From the construction of sets, it follows that $i \notin A'_i$ (otherwise σ' buys some set of cost $\geq \lambda_T/\lambda_i$ and so we may assume $A'_i = A_i$). Consider any scenario-sequence that equals $A_1, \cdots, A_i$ until day i, and has $\{i\}$ as the scenario on day T: the effective cost under σ' is then at least $\lambda_{i+1} \cdot \lambda_T/\lambda_i > \lambda_T$ (since the min-cost set covering i costs λ_T/λ_i and faces inflation at least λ_{i+1}) .
- $A_j = \{j+1, \cdots, T+1\}$ for each $j \leq i$, i.e. σ buys nothing until day i (inclusive). σ' buys nothing before day i and covers $A'_i \neq \emptyset$ on day i. We have the two cases:
  1. If $A'_i \subsetneq A_i$ then let $e \in A_i \setminus A'_i$. Consider any scenario-sequence that equals $A_1, \cdots, A_i$ until day i, and has $\{e\}$ as the scenario on day T: the effective cost under σ' is then at least $2 \cdot \lambda_i > \lambda_T$ (since σ' buys at least two sets on days i or later).
  2. If $A'_i = A_i = \{i+1, \ldots, T+1\}$ then the cost of σ' under any scenario-sequence with $A_1, \cdots, A_i$ until day i is at least $2 \cdot \lambda_i > \lambda_T$ (since the min set cover of $\{i+1, \ldots, T+1\}$ costs at least 2).

In all cases above σ' has objective value strictly larger than λ_T. Thus σ is the unique optimal solution.

Remark: It is still possible that there is always a thrifty (i.e. two-stage) solution to multistage robust set-cover having cost within $O(1) \times \mathsf{Opt}$ (the above example just shows that the constant is at least two). If such a result were true, then in conjunction with the two-stage result from [6] it would also yield Theorem 2.1. However, proving such an existence result seems no easier than obtaining an algorithm for the original multistage problem. Indeed, we take the latter approach in this paper and directly give thrifty approximation algorithms for all the multistage problems considered.

B.2 Bad Example for Subset k-robust Uncertainty sets

Consider the following slight generalization of the k-robust uncertainty model: on each day i we are revealed some set $A_i \in \Omega_i$ such that A_i contains the final scenario A; where $\Omega_i = \{S \subseteq U : |S \cap P_i| \leq k_i\}$ consists of all subsets of U having at most k_i elements in a designated set $P_i \subseteq U$ (it can have any number of elements from $U \setminus P_i$). Again, the final scenario $A = \cap_{i=0}^{T} A_i$.

Recall that we obtain the multistage k-robust model studied in this paper by setting $P_i = U$ (i.e. $\Omega_i = \binom{U}{k_i}$) on all days. We give an example showing that thrifty algorithms perform poorly for multistage set cover under the above 'subset k-robust' uncertainty sets.

Consider a universe U of elements partitioned into T parts $\bigcup_{i=1}^{T} P_i$. We consider a constant number of stages T. The set system consists only of singleton sets, so it suffices to talk about costs on elements. Fix a parameter $\lambda \gg T$. For each $i \in [T]$, we have:

- $|P_i| = \lambda^{i+1}$
- Each P_i-element has cost $1/\lambda^i$
- The inflation factor on day i is λ^i
- $k_i = 1$, so $\Omega_i = \{S \subseteq U : |S \cap P_i| \leq 1\}$

We first show that the optimal value is at most T. Note that on each day i, there is at most one active P_i-element (i.e. $|A_i \cap P_i| \leq 1$). Consider the strategy that on each day i covers the unique active P_i-element: the worst case cost equals $\sum_{i=1}^{T} \lambda^i \cdot 1/\lambda^i = T$. This strategy is feasible since $\bigcup_{i=1}^{T} P_i = U$.

On the other hand, we now show that any strategy that is not active on all days has cost at least $\lambda \gg T$. Consider any strategy that is inactive on some day $j \in [T]$. One the following cases occurs:

1. Suppose that on every scenario sequence, *all* active P_j-elements are covered by day $j-1$. Then consider any scenario sequence with entire P_j active on day $j-1$ (this is possible since $P_j \in \Omega_i$

for all $i \leq j-1$). The total cost incurred on the first $j-1$ days (on this scenario sequence) is at least $|P_j| \cdot \frac{1}{\lambda^j} = \lambda$.

2. Suppose that there is a partial scenario sequence σ (until day $j-1$) where an active P_j-element e remains uncovered after day $j-1$. Extend σ to a scenario sequence that has e active until the end (i.e. e is the unique P_j element that remains active after day j). Since the algorithm is inactive on day j, it must be that e is covered on day $j+1$ or later—thus the total cost along this scenario sequence is at least $\lambda^{j+1} \cdot \frac{1}{\lambda^j} = \lambda$.

Setting $\lambda^{T+1} = \Theta(|U|)$ it follows that thrifty strategies can be worse by a polynomial (in $|U|$) factor for multistage robust set-cover under subset k-robust uncertainty sets.